CST1
COMPUTER SCIENCE TRIPOS Part IB

Thursday 8 June 2023 13:30 to 16:30

COMPUTER SCIENCE Paper 7

*Answer **five** questions.*

*Submit the answers in five **separate** bundles, each with its own cover sheet. On each cover sheet, write the numbers of **all** attempted questions, and circle the number of the question attached.*

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

STATIONERY REQUIREMENTS
Script paper
Blue cover sheets
Tags

SPECIAL REQUIREMENTS
Approved calculator permitted

1 Artificial Intelligence

Evil Robot has decided to toy with his victims by placing them in a puzzle maze. Victims are initially at the `Start` node and need to navigate to the `Finish` node. Other nodes $n_{i,j}$ are laid out on a grid and connected by two kinds of edge. Fixed one-way edges (single-source, solid lines) allow a victim to pass from one node to another, left to right. For example, $n_{5,1}$ to $n_{4,3}$ in the diagram. Switched edges (single-source, solid and dotted lines) can be moved among destinations. For example, $n_{2,3}$ is currently connected to $n_{5,4}$ but can be moved to connect to $n_{1,4}$ or $n_{2,4}$ instead.

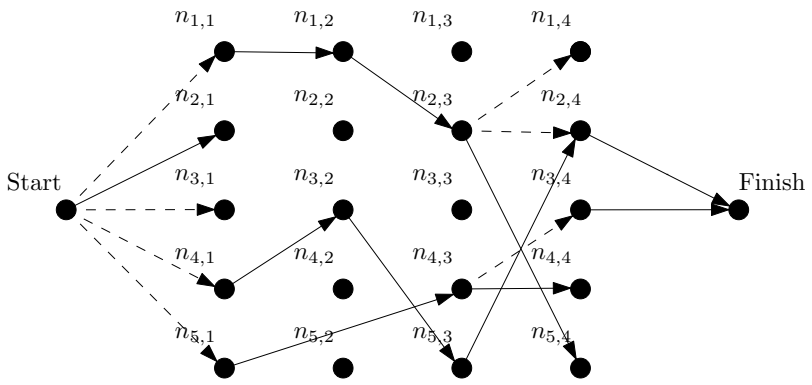

A victim can set the switched edges as desired, but only if they have reached the switch's source node. Their aim is to cross the maze successfully.

(a) The victim decides to solve this problem by representing it as a planning problem using the *state-variable* representation. Give an example of how each of the following elements might appear in the representation.

(i) A *domain*. [1 mark]

(ii) A *rigid relation*. [2 marks]

(iii) A *function*, including a *state variable*, representing the destination of a switched edge. [3 marks]

(iv) A *goal*. [1 mark]

(b) Explain in detail how actions allowing (1) the victim to move around the grid, and (2) the destination of a switched edge to be altered, might be implemented in the state-variable representation. [6 marks]

(c) The victim now wishes to solve their planning problem by converting it to a *constraint satisfaction problem (CSP)*. Explain in detail, giving examples based on the diagram above and your earlier answers, the steps necessary to perform the conversion. [7 marks]

2 Artificial Intelligence

You have a supervised learning problem involving *classification*: a vector $\mathbf{x}$ is to be assigned to one of K classes. To do this you proceed in the usual way: you have a training set $\mathbf{s}$ containing m pairs $(\mathbf{x}_i, \mathbf{y}_i)$. However the labels $\mathbf{y}_i$ are now vectors in $\{0,1\}^K$ containing a single 1 representing the target class. So for example if there are 5 classes and some $\mathbf{x}_i$ should be assigned to class 2 then $\mathbf{y}_i = (0,1,0,0,0)$. To do this, it is proposed that you use K neural networks. The ith network has parameters $\mathbf{w}_i$ and computes the function $h(\mathbf{w}_i, \mathbf{x})$. You may make no further assumptions regarding the function h.

(*a*) You aim to treat the output of the ith network as an estimate of the probability $\Pr(\mathbf{x} \in \text{class } i | \mathbf{x}, \mathbf{w})$ that $\mathbf{x}$ should be in the ith class, where $\mathbf{w}$ collects together all the K vectors $\mathbf{w}_1, \ldots, \mathbf{w}_K$. It is proposed that to do this you should modify the setup described to compute

$$\begin{aligned}\Pr(\mathbf{x} \in \text{class } i | \mathbf{x}, \mathbf{w}) &= \text{prob}(i, \mathbf{x}) \\ &= \frac{\exp(h(\mathbf{w}_i, \mathbf{x}))}{\sum_{j=1}^{K} \exp(h(\mathbf{w}_j, \mathbf{x}))}.\end{aligned}$$

Explain why this modification is required, and how it achieves the stated aim. [4 marks]

(*b*) It is proposed that to train your networks, you should maximize the probability $\Pr(\mathbf{s}|\mathbf{w})$ that a given collection of weights would produce the data in $\mathbf{s}$. (You may consider the training inputs fixed.) Denote by $y_{i,j}$ the jth element of $\mathbf{y}_i$. Show that training can be achieved by minimizing

$$E(\mathbf{w}) = -\sum_{i=1}^{m} \sum_{j=1}^{K} y_{i,j} \log \text{prob}(j, \mathbf{x}_i).$$

State any assumptions that you make. [6 marks]

(*c*) You have previously applied the backpropagation algorithm for training the networks $h(\mathbf{w}_i, \mathbf{x})$ and as a result of this you know how to compute derivatives $\partial h(\mathbf{w}_i, \mathbf{x})/\partial w_{i,j}$ where $w_{i,j}$ is the jth element of $\mathbf{w}_i$. Explain what further steps are necessary to use this knowledge to obtain derivatives of $E(\mathbf{w})$ with respect to the relevant weights. [10 marks]

3 Economics, Law and Ethics

(*a*) Describe the six branches of philosophy, explaining what insight they might give into technology ethics. [12 marks]

(*b*) Using a consequentialist approach to the philosophy of ethics, outline the ethical implications of facial recognition technology. Your answer should identify the potential benefits and harms of using the technology. [8 marks]

4 Economics, Law and Ethics

(*a*) How does game theory account for the distribution of trolls (accounts that make deliberately offensive and provocative posts) on social media platforms? [10 marks]

(*b*) If a popular social media platform changes its moderation policy, allowing trolls to proliferate, when are users likely to switch to another platform? [10 marks]

5 Formal Models of Language

The Earley parser is an example of a top-down LR chart parser. This question requires you to design a *bottom-up RL chart parser* for parsing sentences with a context-free grammar.

(*a*) Describe the structure of the chart. [4 marks]

(*b*) Explain how the chart is initialised. [2 marks]

(*c*) Describe the parsing algorithm using rule induction notation to describe new edges that can be added to the chart. [6 marks]

(*d*) Describe the properties of the chart when a parse is found. [2 marks]

(*e*) Explain how to reconstruct a parse tree from the chart. [2 marks]

(*f*) Explain how the algorithm could choose between parses if more than one is found. [2 marks]

(*g*) Does your approach result in a compact chart? If not, how could you reduce the number of edges in a chart without losing information about the possible parses? [2 marks]

6 Formal Models of Language

Consider the following syntax tree:

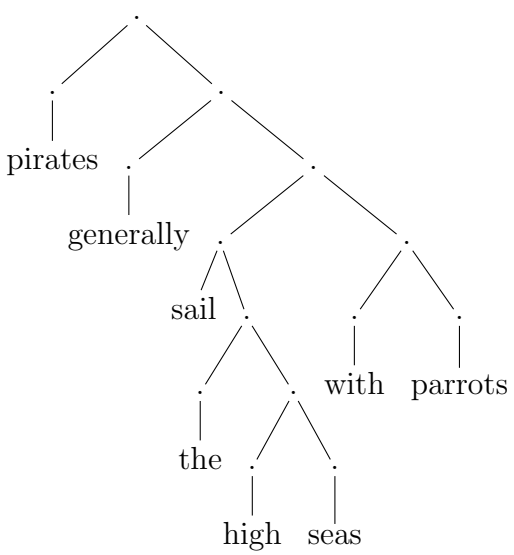

(*a*) Provide a Context-Free Grammar that could generate the tree. [5 marks]

(*b*) Provide a Tree-Adjoining Grammar that could generate the tree. [5 marks]

(*c*) Show how to generate the tree using your Tree-Adjoining Grammar. [5 marks]

(*d*) Could the following Grammatical Relations be derived from the tree? Explain your answer.

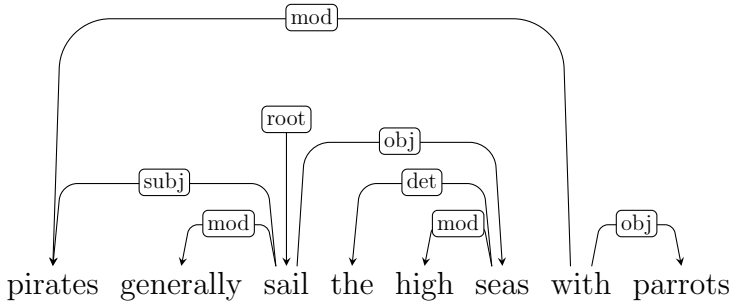

[5 marks]

7 Further Graphics

(a) State the best geometry representation for the following tasks. Briefly explain.

(i) Testing if a line segment intersects a surface. [1 mark]

(ii) Tracking the surface of a fluid undergoing geometry and topology (connectivity) changes. [1 mark]

(iii) Aligning and merging surfaces acquired from different views. [1 mark]

(iv) Rendering the surface of a cube. [1 mark]

(v) Look up the colour of each point on a surface from a texture. [1 mark]

(b) Assume constant minimum κ_{min} and maximum κ_{max} curvature for a surface. Which of the following could be a closed surface?

(i) $\kappa_{min} = \kappa_{max}$. [1 mark]

(ii) $2H^2 - K = 0$ for mean curvature H and Gaussian curvature K. [2 marks]

(c) A heightfield is a surface in 3D defined by a function $h(x, y)$ over the xy-plane. Assuming we define points with $z > h(x, y)$ as outside this surface, provide an expression for the surface normal in terms of the derivatives of h without using the cross product. [4 marks]

(d) Given a surface in both parametric $\mathbf{p}(u, v)$ and implicit $f(\mathbf{p}) = 0$ forms,

(i) prove that the surface normal of the parametric form is parallel to the normal of the implicit form.
[*Hint:* Chain rule: $\frac{\partial f(a(x),b(x),c(x))}{\partial x} = \frac{\partial f}{\partial a}\frac{\partial a}{\partial x} + \frac{\partial f}{\partial b}\frac{\partial b}{\partial x} + \frac{\partial f}{\partial c}\frac{\partial c}{\partial x}$] [6 marks]

(ii) Will the normals remain parallel if the implicit function is also a signed distance function? [1 mark]

(iii) Why is the sign ambiguous? [1 mark]

8 Further Graphics

(a) We are given a scene with a point light source at a point $\mathbf{p}$ that emits light of radiance L in all directions, and a plane with equation $\mathbf{n}^T(\mathbf{x}-\mathbf{c}) = 0$ and BRDF f_r. We send a ray from the camera sensor with origin $\mathbf{o}$ and direction $\mathbf{d}$ such that each point on the ray is $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ for $t > 0$.

(i) What is the intersection point of the ray and the plane? [2 marks]

(ii) What is the incoming light direction? [1 mark]

(iii) What is the direction of the outgoing reflected light towards the sensor origin $\mathbf{o}$? [1 mark]

(iv) What is the radiance of the light reflected towards the sensor origin $\mathbf{o}$? [2 marks]

(b) A surface represented by an implicit function $f(\mathbf{x})$ is transformed by linear blend skinning such that each point $\mathbf{x}$ is mapped to $\mathbf{x}'$. There are two bones with transformation matrices $\mathbf{T}_1$ and $\mathbf{T}_2$ and corresponding blending weights $w_1(\mathbf{x})$ and $w_2(\mathbf{x}) = 1 - w_1(\mathbf{x})$. In addition, we are given $s_1(\mathbf{x}') = w_1(\mathbf{x})$ and $s_2(\mathbf{x}') = w_2(\mathbf{x})$.

(i) Write the expression for $\mathbf{x}'$. [1 mark]

(ii) Simplify the expression for $\mathbf{x}'$ assuming no rotations are stored on the bones. Modify the expression further for the case where the weights are not summing up to 1. [2 marks]

(iii) What is the implicit function of the transformed surface assuming all blended rotations are invertible? Simplify as much as possible. [3 marks]

(iv) What is the implicit function of the transformed surface if there are no rotations stored on the bones? Simplify as much as possible. [2 marks]

(c) In this question, we will think about how quaternions represent spatial rotations.

(i) Write the form of a general quaternion and the form of a quaternion representing rotations. [2 marks]

(ii) Prove that a quaternion $\mathbf{q}$ and $-\mathbf{q}$ represents the same rotation. [2 marks]

(iii) Given a quaternion $\mathbf{q}$ with no real part, what is $\mathbf{qq}$? [2 marks]

9 Further Human–Computer Interaction

British houses are often wastefully heated even when there is nobody home. The problem is that it takes a long time to heat up to a comfortable temperature, and it is too difficult to program thermostats to account for the variable routines an occupant might have from one day to the next.

Imagine someone who owns an Internet-connected thermostat, a semi-autonomous car, and a voice-activated smart speaker. A generative language model fine-tuned for dialog, like the recently popular ChatGPT, could be used via the smart speaker to define a policy such that, on cold days, the heating will come on when the car will be returning home in about 30 minutes.

(*a*) Complete the following voice dialog that might be used by this person to define the above heating policy for the next year.

```
USER: I want to define a new policy to control the
  central heating
BOT: What is the new policy?
USER: If the outside temperature is less than 10 degrees,
  use a special winter policy
BOT: What is the special winter policy?
```

[2 marks]

(*b*) Which wave of HCI theory is most relevant to this interaction scenario, and why? [2 marks]

(*c*) Suggest *four* cognitive dimensions that are relevant to the dialog. [4 marks]

(*d*) For *each* dimension, explain:

(*i*) How this dimension is relevant to the voice dialog.

(*ii*) What its significance is in relation to the user's goals.

(*iii*) How a voice dialog system might behave differently to improve usability on this dimension.

[12 marks]

10 Further Human–Computer Interaction

In future, many households will have a limited energy budget. Diverse household members are likely to have different energy requirements that must be negotiated within that fixed overall budget.

This question relates to the design of a shared display that can be used by household members to monitor and discuss their respective energy usage. The display will be primarily graphical, rather than using text or numbers.

(*a*) Make a sketch of the display layout, which addresses the scenario above, and can be used to answer the remainder of the question. Use arrows to label the four elements referred to in the next part. [2 marks]

(*b*) Name *four* different types of graphical mark or property that have been used as the basic elements of the overall design. [2 marks]

(*c*) For *each* of these types of graphical mark, explain:

(*i*) What it looks like (its appearance).

(*ii*) What it represents (its meaning).

(*iii*) The nature of the correspondence between its appearance and meaning in your design.

[12 marks]

(*d*) Describe how qualitative empirical data could be collected and analysed to evaluate and improve this design concept. [4 marks]

END OF PAPER